



# UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

Address: COMMISSIONER FOR PATENTS

P.O. Box 1450

Alexandria, Virginia 22313-1450

www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/612,613	07/01/2003	Douglas B. Davis	RSW920030040US1 (094)	3998
46320 7590 02/27/2009 CAREY, RODRIGUEZ, GREENBERG & PAUL, LLP STEVEN M. GREENBERG 950 PENINSULA CORPORATE CIRCLE SUITE 3020 BOCA RATON, FL 33487				
EXAMINER MEHRMANESH, ELMIRA				
ART UNIT 2113		PAPER NUMBER		
MAIL DATE 02/27/2009		DELIVERY MODE PAPER		

**Please find below and/or attached an Office communication concerning this application or proceeding.**

The time period for reply, if any, is set in the attached communication.

UNITED STATES PATENT AND TRADEMARK OFFICE

---

BEFORE THE BOARD OF PATENT APPEALS  
AND INTERFERENCES

---

*Ex parte* DOUGLAS B. DAVIS, YIH-SHIN TAN,  
BRAD B. TOPOL, and VIVEKANAND VELLANKI

---

Appeal 2008-2281  
Application 10/612,613  
Technology Center 2100

---

Decided:<sup>1</sup> February 27, 2008

---

Before JOSEPH L. DIXON, ALLEN R. MACDONALD,  
and JEAN R. HOMERE, *Administrative Patent Judges*.

MACDONALD, *Administrative Patent Judge*.

DECISION ON APPEAL

---

<sup>1</sup> The two-month time period for filing an appeal or commencing a civil action, as recited in 37 CFR § 1.304, begins to run from the decided date shown on this page of the decision. The time period does not run from the Mail Date (paper delivery) or Notification Data (electronic delivery).

Appellants appeal under 35 U.S.C. § 134 from the Examiner's rejection of claims 1-24. We have jurisdiction under 35 U.S.C. § 6(b).

We reverse.

These opinions follow: (1) a majority opinion authored by Judge MacDonald, joined by Judge Dixon, and (2) a dissenting opinion authored by Judge Homere. Judge Homere would affirm the Examiner's rejection of all claims.

MacDONALD, *Administrative Patent Judge*.

#### STATEMENT OF THE CASE

According to Appellants the invention relates to checkpointing technology where long running Web services and their respective execution states can be revived in response to an event.<sup>2</sup>

Claims 1 and 3 are illustrative:

1. A checkpoint processor configured for coupling to individual Web services through a Web services engine, said checkpoint processor comprising:

checkpoint logic programmed to store checkpoint data for the individual Web service instance invocations;

restart logic programmed to restore said stored checkpoint data to a replacement for failed ones of the individual Web service instance invocations; and,

---

<sup>2</sup> See Spec. ¶9.

cleanup logic programmed to removed said stored checkpoint data for concluded, non-failed ones of the individual Web service instance invocations.

3. A method for managing checkpoints in a Web application, the method comprising the steps of:

storing a state object for an invocation of a requesting Web service instance; and,

responsive to a failure in said Web service instance, restarting a replacement Web service instance and providing said state object to a replacement Web service instance for said requesting Web service instance.

The Examiner relies on the following prior art reference:

Doyle

US 2004/0243915 A1

Dec. 2, 2004

The Examiner rejected claims 1-24 under 35 U.S.C. § 102(e) as anticipated over Doyle (Ans. 3-10).

*Examiner's Findings and Conclusions*

The Examiner defined “checkpoint” as a “file containing information that describes the state of the system (environment) at a particular time” (Ans. 10). The Examiner found the “Optimization Metrics 180” of Doyle “can include a listing of various static and dynamic parameters and measurements associated with the operation of the individual service instances 130A, 130B” (Ans. 10). The Examiner concluded that the

“‘Optimization Metrics’ are identical to the claimed ‘checkpoint data’ as they contain information that describes the state of the system/environment (i.e., individual service instances 130A, 130B) at a particular time” (Ans. 11).

The Examiner also concluded that Doyle teaches restoring checkpoint data to a replacement for failed ones of the individual Web service instance invocations (Ans. 11). The Examiner found that Doyle discloses (1) detecting a node failure within the grid coordinator; (2) identifying each service instance residing within the failed node; (3) identifying one or more replacement nodes; (4) retrieving the logged metrics for each of the service instances in the failed node; and (5) identifying platform metrics for each of the replacement nodes (Ans. 11).

The Examiner also found that Doyle discloses that the “optimization logic 220 initially can determine whether services 260A, 260C hosted within the failed node 210Y can be wholly reinstantiated in one of the identified replacement nodes 240A, 240B in a manner in which the services 260A, 260C can continue to operate at a performance level previously attained in the failed node 210Y” (emphasis omitted) (Ans. 12).

Further, the Examiner found that “Doyle’s failover process includes detecting the faulty node and providing continuous operation of the failed instances by retrieving the checkpoint data (i.e., logged metrics) of the failed service instances to enable their full resumption in the replacement nodes based upon the collected metrics” (Ans. 12). Thus, the Examiner concluded that “[i]n order to maintain continuous operations of the failed service instances and to fully re-instantiate the failed service instances, it is essential to provide the previously stored checkpoint data in the replacement

invocations (i.e., restore the checkpoint data) to enable restarting of the service instances from the previously stored state of their operations prior to the failure” (emphasis omitted) (Ans. 12). The Examiner then concluded that “Doyle discloses restoring checkpoint data to a replacement for failed ones of the service instance invocations” (Ans. 12).

*Appellants' Contentions*

Appellants argue that the Doyle fails to teach “restoring checkpoint data to a replacement for failed ones of the individual Web service instance invocations, as claimed” (underlining omitted) (App. Br. 11). Appellants admit that Doyle teaches that “the logged metrics for each of the service instances in a failed node are retrieved” (App. Br. 11). Further, Appellants admit that Doyle teaches “using service metrics 230 and platform metrics 250A, 250B in a best-fit analysis” that “is used to determine the replacement node in which new instances of the services are created” (App. Br. 11).

Appellants also argue that the “service metrics 230 are used to place new instances of the services across one or more replacement nodes in an optimal fashion,” and does not disclose the claimed restoring said checkpoint data to a replacement (App. Br. 8). Additionally, Appellants argue that “[u]sing service metrics to determine a replacement node ... is not identical to restoring checkpoint data to a replacement” (App. Br. 10). Further, Appellants argue that the “optimization/service metrics are used to select the replacement nodes” (underlining omitted) (Reply Br. 3).

## ISSUE

The issue before us, then, is whether Doyle teaches restoring the stored checkpoint data to a replacement for failed ones of the individual Web service invocations.

## FINDINGS OF FACT

1. Claim 1 requires “restart logic programmed to restore said stored checkpoint data to a replacement for failed ones of the individual Web service instance invocations” (App. Br. 16).
2. Doyle discloses that “[u]pon detecting the failure, a set of replacement nodes can be located and platform metrics can be determined for each of the replacement nodes” (Doyle ¶12). Doyle also discloses that “[c]onsequently new replacement service instances can be created in optimally selected ones of the replacement nodes based upon the collected metrics and the platform metrics” (*Id.*). Further, Doyle discloses “[i]n this regard, a best-fit analysis can be performed between the collected metrics and the platform metrics to identify an optimal replacement node in which to create new replacement service instances” (*Id.*)
3. Doyle discloses that “optimization metrics 180 can include a measurement of the resources consumed in a grid host.” The optimization metrics can also “specify for each individual service instance 130A, 130B a cost per unit of performance, a cost per unit of resource consumption, revenue per unit of performance, and revenue per unit of consumption.” “[T]he data included in the store of optimization metrics 180 can be updated regularly by operation of a monitor (not shown) coupled to the store of

optimization metrics 180 which can collect performance data for the individual service instances 130A, 130B" (*See* Doyle ¶36).

4. Doyle discloses that "where replacement node 240A has been determined to include computing resources able to accommodate new instances of the services 260A, 260C at a level of performance previously attained in the failed node 210Y, new instances of the services 260A, 260C can be placed within the replacement node 240A" (Doyle ¶40).

5. Doyle discloses that block 340 teaches logged metrics for each of the service instances in the failed node can be retrieved so as to smooth anomalous events in any one instance of the service, and in any case, "platform metrics for each of the identified replacement nodes can be identified" (Doyle ¶43). Doyle also discloses that if replacement nodes, which can accommodate new instances of the services without requiring a reduction in performance, are identified, then new instances of the services, can be placed within the identified replacement node without a performance reduction (*See* Doyle ¶44).

6. The Specification of the present application discloses that when restarting a Web service instant invocation is required, the restarted Web service instance invocation can be restored to its former state by uploading each of the instance identifier, the persisted state information, and the asynchronous correlator (Spec. ¶19). The Specification of the present application also discloses that "the checkpoint processor 160 can receive and persist state objects 180 provided by the instances of Web services 130A, 130B, 130n in a persistent store 170" (Spec. ¶21).



## PRINCIPLES OF LAW

In rejecting claims under 35 U.S.C. § 102, “[a] single prior art reference that discloses, either expressly or inherently, each limitation of a claim invalidates that claim by anticipation.” *Perricone v. Medicis Pharmaceutical Corp.*, 432 F.3d 1368, 1375 (Fed. Cir. 2005), citing *Minn. Mining & Mfg. Co. v. Johnson & Johnson Orthopaedics, Inc.*, 976 F.2d 1559, 1565 (Fed. Cir. 1992). “Anticipation of a patent claim requires a finding that the claim at issue ‘reads on’ a prior art reference.” *Atlas Powder Co. v. IRECO, Inc.*, 190 F.3d 1342, 1346 (Fed. Cir. 1999) (“In other words, if granting patent protection on the disputed claim would allow the patentee to exclude the public from practicing the prior art, then that claim is anticipated, regardless of whether it also covers subject matter not in the prior art.”) (Internal citations omitted).

“[A]nticipation of a claim under § 102 can be found only if the prior art reference discloses every element of the claim . . . .” *In re King*, 801 F.2d 1324, 1326 (Fed. Cir. 1986) (citing *Lindemann Maschinenfabrik GMBH v. American Hoist & Derrick Co.*, 730 F.2d 1452, 1457 (Fed. Cir. 1984)). “[A]bsence from the reference of any claimed element negates anticipation.” *Kloster Speedsteel AB v. Crucible, Inc.*, 793 F.2d 1565, 1571 (Fed. Cir. 1986).

“A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference.” *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631 (Fed. Cir. 1987). Analysis of whether a claim is patentable over the prior art under 35 U.S.C. § 102 begins with a determination of the scope of the claim. We determine the scope of the claims in patent applications

not solely on the basis of the claim language, but upon giving claims their broadest reasonable construction in light of the specification as it would be interpreted by one of ordinary skill in the art. *In re Am. Acad. of Sci. Tech. Ctr.*, 367 F.3d 1359, 1364 (Fed. Cir. 2004). The properly interpreted claim must then be compared with the prior art.

### ANALYSIS

We now consider the Examiner's anticipation rejection of claims 1-24 over Doyle. Claim 1 requires restoring the stored checkpoint data to a replacement for failed ones of the individual Web service instance invocations (FF 1).

We find that Doyle does not expressly teach providing the functionality to perform restoring said restored checkpoint data to a replacement for failed ones of Web services instances. We acknowledge that Doyle teaches:

- (1) retrieving logged metrics;
  - (2) utilizing the logged metrics, for each of the service instances, to smooth anomalous events in any one instance of the service; and
  - (3) identifying platform metrics for the replacement nodes (FF 5).
- Doyle also discloses using "collected metrics" and "platform metrics" to identify an optimal replacement node (FF 2).

The Examiner concluded that the "'Optimization Metrics' are identical to the claimed 'checkpoint data' as they contain information that describes the state of the system/environment (i.e., individual service instances 130A, 130B) at a particular time" (Ans. 11). Thus, we find that one of ordinary skill in the art would have recognized that the "logged

metrics," "collected metrics," and "platform metrics" are Optimization Metrics. However, while Doyle discloses utilizing the "logged metrics," "collected metrics," and "platform metrics," Doyle fails to expressly teach that the "logged metrics," "collected metrics," or "platform metrics" consists of data that is restored to a replacement for failed instances of Web services.

Doyle also discloses that where a replacement node has been identified that can accommodate new instances of services at a level previously attained, the new instances of the services can be placed within the replacement node (FF 4). However, Doyle does not expressly teach that the "logged metrics," "collected metrics," or "platform metrics are restored to the new instance of services at the replacement node.

Doyle also clearly teaches using metrics to identify a replacement node that can accommodate the failed instances of services (FF 2). However, Doyle does not disclose restoring stored "logged metrics," "collected metrics," or "platform metrics" (*i.e.*, checkpoint data) to a replacement for the failed ones of the individual Web service instances, as required by claim 1.

Therefore, for the foregoing reasons, Appellants have persuaded us of error in the Examiner's rejection of independent claim 1. Therefore, we will reverse the Examiner's rejection of that claim, and claim 2 which falls with claim 1.

Independent claims 3 and 14 are commensurate in scope with independent claim 1, as claims 3 and 14 require "storing a state object" and "responsive to a failure in said Web service instance," "providing said state object to a replacement Web service instance." The Specification of the present application does not expressly disclose a definition of "state object."

However, we find that one of ordinary skill in the art would have recognized that the state object was equivalent to "checkpoint data," as (1) the checkpoint processor could receive and persist state objects provided by the instances of Web services in a persistent store (FF 6), and (2) a restarted Web service instance invocation could be restored to its former state by uploading the persisted state information (FF 6). Thus, one of ordinary skill in the art would have characterized the persisted state information (*i.e.*, persisted state objects) as checkpoint data, as it was the checkpoint processor that received and persisted the state objects.

Thus, for the reasons discussed above with regards to independent claim 1, we will reverse the Examiner's rejection of (1) independent claim 3 and claims 4-13 that fall with that claim, and (2) independent claim 14, and claims 15-24 that fall with that claim.

#### DECISION

We have reversed the Examiner's rejections with respect to all claims on appeal. Therefore, the Examiner's decision rejecting claims 1-24 is reversed.

#### REVERSED

HOMERE, *Administrative Patent Judge, dissenting.*

I write separately to voice my disagreement with the majority's finding that Appellants have shown that Doyle does not teach restoring a restored checkpoint data to a replacement for failed ones of Web services instances, as allegedly recited independent claim 1. Because of this finding, the majority reverses the Examiner's prior art rejection of claims 1 through 24. From that decision, I respectfully dissent.

The majority opinion states *inter alia*:

Doyle also clearly teaches using metrics to identify a replacement node that can accommodate the failed instances of services (FF 2). However, Doyle does not disclose restoring stored "logged metrics," "collected metrics," or "platform metrics" (*i.e.*, checkpoint data) to a replacement for the failed ones of the individual Web service instances, as required by claim 1.

Therefore, for the foregoing reasons, Appellants have persuaded us of error in the Examiner's rejection of independent claim 1.

(Majority Op. 10).

I note in the outset that claim 1 recites in relevant part a restart logic *programmed to restore* a stored checkpoint data to a replacement for failed ones of the individual Web service instance invocations. The majority

disregards<sup>3</sup> the claim recitation *programmed to*, and instead finds that the claim requires a restart logic that restores a stored checkpoint data. I do not agree. In my view, such a narrow interpretation of the claim language is not consistent with *In re Bigio*, which requires that the claim be given the broadest reasonable interpretation. *In re Bigio*, 381 F.3d 1320, at 1324 (Fed. Cir. 2004). The cited claim language merely requires the restart logic to be *programmed to restore* the stored checkpoint data. Such a recitation does not require the restart logic to *actually restore* the checkpoint data. At best, the recited limitation requires that the restart logic be *capable of* restoring the stored checkpoint data in the event of a failure. Therefore, the recitation is simply a statement of intended use, which should be afforded no patentable weight. See MPEP § 2114. A statement of intended use in an apparatus claim cannot be used to distinguish the claim over the prior art apparatus. See *In re Schreiber*, 128 F.3d 1473, 1477 (Fed. Cir. 1997). A claimed apparatus must be described by its structure, not its intended use. The mere recitation of an intended use in a claim will not be given any patentable

---

<sup>3</sup> When presented with a claim including nonfunctional descriptive material, an Examiner must determine whether such material should be given patentable weight. The Patent and Trademark Office (PTO) must consider all claim limitations when determining patentability of an invention over the prior art. *In re Gulack*, 703 F.2d 1381, 1385 (Fed. Cir. 1983). The PTO may not disregard claim limitations comprised of printed matter. See *Gulack*, 703 F.2d at 1384; see also *Diamond v. Diehr*, 450 U.S. at 191. However, the PTO need not give patentable weight to descriptive material absent a new and unobvious functional relationship between the descriptive material and the substrate. See *Gulack*, 703 F.2d at 1386. See also *In re Ngai*, 367 F.3d 1336, 1338 (Fed. Cir. 2004); *In re Lowry*, 32 F.3d 1579, 1583-84 (Fed. Cir. 1994). The burden of establishing the absence of a novel, nonobvious functional relationship rests with the PTO. *In re Lowry*, 32 F.3d at 1584.

weight. *Application of Dense*, 156 F.2d 76 (CCPA 1946). *See also Ex Parte James A. Satchell*, Appeal 2008-0071, 2008 WL 4828136, (BPAI 2008). Further, because the restart logic is only programmed to restore a stored checkpoint data, as opposed to actually restoring the stored checkpoint data, such data is not performing any function in the claim. It is sought to be distinguished for its nature, and not for its function. In my view, the checkpoint data is also non-functional descriptive material, which should not be given any patentable weight.<sup>4</sup>

Consequently, Appellants' argument that Doyle does not teach restoring the stored checkpoint data is not commensurate in scope with claim 1. Therefore, it is not persuasive. I am thus satisfied that Appellants have not shown that the Examiner erred in finding that Doyle anticipates claim 1.

Thus, I cannot agree with the majority's reversal of the Examiner's rejection of claim 1. Since the majority's reversal of the Examiner's rejection of claims 2 through 24 is hinged upon the same rationale provided for reversing claim 1, and not on the independent merits of the cited claims, I take exception to the reversal of claims 2 through 24 for the same reasons articulated above. Consequently, I would affirm the Examiner's rejection of claims 1 through 24.

---

<sup>4</sup> *Id.*

Appeal 2008-2281  
Application 10/612,613

pgc

CAREY, RODRIGUEZ, GREENBERG & PAUL, LLP  
STEVEN M. GREENBERG  
950 PENINSULA CORPORATE CIRCLE  
SUITE 3020  
BOCA RATON, FL 33487